

# MASSIVELY PARALLEL-IN-SPACE-TIME, ADAPTIVE FINITE ELEMENT FRAMEWORK FOR NON-LINEAR PARABOLIC EQUATIONS

ROBERT DYJA, BASKAR GANAPATHYSUBRAMANIAN,  
AND KRISTOFFER G. VAN DER ZEE

**ABSTRACT.** We present an adaptive methodology for the solution of (linear and) non-linear time dependent problems that is especially tailored for massively parallel computations. The basic concept is to solve for large blocks of space-time unknowns instead of marching sequentially in time. The methodology is a combination of a computationally efficient implementation of a parallel-in-space-time finite element solver coupled with *a posteriori* space-time error estimates and a parallel mesh generator. This methodology enables, in principle, simultaneous adaptivity in both space and time (within the block) domains. We explore this basic concept in the context of a variety of time-steppers including  $\Theta$ -schemes and Backward Differentiate Formulas. We specifically illustrate this framework with applications involving time dependent linear, quasi-linear and semi-linear diffusion equations. We focus on investigating how the coupled space-time refinement indicators for this class of problems affect spatial adaptivity. Finally, we show good scaling behavior up to 150,000 processors on the Blue Waters machine. This is achieved by careful usage of memory via block storage and non-zero formats along with lumped communication for matrix assembly. This methodology enables scaling on next generation multi-core machines by simultaneously solving for large number of time-steps, and reduces computational overhead by refining spatial blocks that can track localized features. This methodology also opens up the possibility of efficiently incorporating adjoint equations for error estimators and inverse design problems, since blocks of space-time are simultaneously solved and stored in memory.

## 1. INTRODUCTION

We describe the methodology, implementation details and application examples of space-time block adaptive solutions to parabolic partial differential equations (evolution problems). This approach is primarily motivated by the necessity of

---

CZESTOCHOWA UNIVERSITY OF TECHNOLOGY, UL. DABROWSKIEGO 73, 42-201 CZESTOCHOWA, POLAND

IOWA STATE UNIVERSITY, 306 LAB OF MECHANICS, AMES, IA 50011

UNIVERSITY OF NOTTINGHAM, SCHOOL OF MATHEMATICAL SCIENCES, UNIVERSITY PARK, NOTTINGHAM NG7 2RD, UK

*E-mail addresses:* robert.dyja@icis.pcz.pl, baskarg@iastate.edu,  
KG.vanderZee@nottingham.ac.uk.

*Key words and phrases.* parabolic problems, parallel-in-time, finite element method, adaptive mesh refinement.

We gratefully acknowledge NSF 1435587 and XSEDE resources at TACC, as well as an exploratory account on BlueWaters.

Kristoffer G. van der Zee's research was supported by the Engineering and Physical Sciences Research Council (EPSRC) under grant EP/I036427/1.

designing computational methodologies that can scale to leverage the availability of very large computing clusters (with  $\sim 10$ 's to  $100$ 's of thousands of processors). For evolution problems, the standard approach of decomposing the spatial domain is a powerful paradigm of parallelization. However, for a fixed spatial discretization, the efficiency of purely spatial domain decomposition degrades substantially beyond a threshold (usually  $\sim$  thousands of processors, beyond which efficiency is communication limited) which make this approach unsuitable on larger machines. To overcome this barrier, a natural approach is to consider the time domain as an additional dimension and simultaneously solve for blocks of time, instead of the standard approach of sequential time-stepping. Early work on this approach was considered by Hughes and coworkers [15, 16], Tezduyar and co-workers [24], and Reddy and co-workers [21], while variations on this theme have recently been explored by several groups [6, 9, 18, 19, 22, 28]. Other examples include methods that exploit the fact that some coefficients during matrix assembly can be calculated concurrently [7] and methods that transform the equations so as to build one system of equations for more than one timestep [17].

The concept of solving for blocks of time simultaneously has recently gained a lot of attention to enable effective usage of exascale computing resources (see for instance the US DOE's Exascale Mathematics Working Group [14]). In addition to this obvious advantage, solving for space-time blocks also allows natural incorporation of *a posteriori* error estimates for mesh adaptivity, and enables the solution of inverse problems (involving adjoints). This has several additional tangible benefits in the context of computational overhead. For evolution problems that exhibit regionalized behavior in space and time (wave equations, moving interfaces like bubbles and shocks) solving in blocks of space-time that are locally refined to match the regional behavior provides substantial computational gain [9]. Similarly, the availability of error estimates across a block of time allows optimal choices of space and time adaptivity.

Motivated by these considerations, this paper presents a methodology for the solution of linear, quasi-linear and semi-linear (time dependent) diffusion equations in three dimensions. We discuss the development of a parallel adaptive framework for the solution of large blocks of space-time. We detail the development of the block space-time framework for two classes of time-steppers ( $\theta$  schemes, Backward Difference Formula (BDF)). We discuss implementation details of the massively parallel adaptive block-space time framework and illustrate scaling behavior up to 150,000 processors. We subsequently define *a posteriori space-time error indicators* to identify spatial regions for mesh adaption. Finally, we demonstrate that for sufficiently large problems the block space-time approach is much more computationally efficient than the standard sequential time-stepping approach.

The outline of the rest of the paper is as follows: Section 2, and 3 detail the block space-time framework for linear, and non-linear evolution equations, respectively. Section 4 discusses the space-time error estimates for these classes of problems. In Section 5, we discuss implementation details and show scaling performance and analysis. Section 6 illustrates several numerical examples of the framework. We conclude in Section 7.

## 2. BASIC SPACE-TIME FORMULATION: LINEAR AND NON-LINEAR VERSIONS

**2.1. Space-time framework for a linear problem.** Given a bounded domain  $\Omega \in \mathbb{R}^3$ , and a finite time domain  $[0, T]$ , consider the parabolic equation that solves for  $u: \Omega \times [0, T] \rightarrow \mathbb{R}$ :

$$(1) \quad \begin{cases} \partial_t u(\mathbf{x}, t) - \nabla \cdot \kappa \nabla u(\mathbf{x}, t) = f(\mathbf{x}, t) & \text{in } \Omega \times [0, T] \\ u(\mathbf{x}, 0) = u_0 \end{cases}$$

where  $f: \Omega \times [0, T] \rightarrow \mathbb{R}$  is a smooth source function, and  $\kappa > 0$ . We consider (without loss of generality) that Dirichlet boundary conditions are imposed on the boundary  $\Gamma$ , unless otherwise specified. Considering a tessellation,  $\mathcal{T} \equiv \{\Omega^1, \dots, \Omega^e, \dots\}$ , of the domain  $\Omega$  into elements (with average size  $h$ ), the weak form of this equation is given as:

$$(2) \quad \begin{cases} \text{Find } u^h(\cdot, t) \in \mathcal{U}^h : \\ (w^h, \partial_t u^h(\cdot, t)) + (\nabla w^h, \kappa \nabla u^h(\cdot, t)) = (w^h, f) \quad \forall w^h \in \mathcal{V}^h \end{cases}$$

where  $(\cdot, \cdot)$  is the  $L_2$  inner product on  $\Omega$  and

$$(3) \quad \begin{aligned} \mathcal{U}^h &:= \{u^h | u^h \in H^1(\Omega), \quad u^h \in P(\Omega^e) \quad \forall e\} \\ \mathcal{V}^h &:= \{w^h | w^h \in H^1(\Omega), \quad w^h \in P(\Omega^e) \quad \forall e\} \end{aligned}$$

with  $P(\Omega^e)$  being the space of the standard polynomial finite element shape functions on element  $\Omega^e$ . To obtain a fully discretized form, we employ a time stepping technique on the above semi-discrete equation. While any time-stepping method can be used, as an example, consider the Euler Backward Formula that is defined on a discretization  $\{0, t_1, t_2, \dots, T\}$  of the time domain:

$$(4) \quad \left( w^h, \frac{u_{n+1}^h - u_n^h}{\Delta t} \right) - (\nabla w^h, \kappa \nabla u_{n+1}^h) = (w^h, f_{n+1}), \text{ for } n = 0, 1, \dots$$

where the subscript denotes evaluation at that discrete time, and  $\Delta t = t_{n+1} - t_n$  is the timestep.

Following standard FEM practice, with the tessellation of the domain resulting in  $k$  nodal values (to describe spatial variation) of the field  $u$ , equation (4) can be expressed in terms of matrix-vector products as:

$$(5) \quad \mathbf{M} \mathbf{u}_{n+1} - \Delta t \mathbf{K} \mathbf{u}_{n+1} = \mathbf{M} \mathbf{u}_n + \Delta t \mathbf{f}_{n+1}, \text{ for } n = 0, 1, \dots$$

where  $\mathbf{M}$  and  $\mathbf{K}$  are the global mass and stiffness matrices, respectively.<sup>1</sup>  $\mathbf{u}_{n+1}$  and  $\mathbf{u}_n$  are vectors containing the nodal values of the field  $u$  at time step  $n+1$  and  $n$ , respectively. Equation (5) represents the system of equations solved to get the solution for timestep  $n+1$ . The size of vector  $\mathbf{u}$  is equal to the number of nodal unknowns,  $k$ . Similarly matrices  $\mathbf{K}$ ,  $\mathbf{M}$  are sparse matrices of size  $k \times k$ .

Consider a block-wise division of the total time domain. Each block,  $B_i$ , consists of multiple timesteps. This is schematically represented in Figure 1. Instead of sequentially solving for each time step (as in (5)), consider solving for the field variable in a complete time block,  $B$ , consisting of  $N$  timesteps simultaneously, i.e. solve for  $\mathbf{u}_i$ ,  $i = 1, \dots, N$  simultaneously.

<sup>1</sup>With some abuse of notation, the elements of matrix  $\mathbf{M}$  are equal to  $M_{ij} = (w_i^h, w_j^h)$  and matrix  $\mathbf{K}$  are equal to  $K_{ij} = (\nabla w_i^h, \kappa \nabla w_j^h)$ .




FIGURE 1. Indexing of timesteps per block (B). Number of timesteps per block is equal to N.

This results in a block diagonal matrix of size  $(N \times k)$ -by- $(N \times k)$  given by:

$$(6) \quad \begin{bmatrix} \mathbf{I} & & & & \\ -\mathbf{M} & \mathbf{M} + \Delta t \mathbf{K} & & & \\ & -\mathbf{M} & \mathbf{M} + \Delta t \mathbf{K} & & \\ & & \ddots & \ddots & \\ & & & -\mathbf{M} & \mathbf{M} + \Delta t \mathbf{K} \end{bmatrix} \begin{bmatrix} \mathbf{u}_0 \\ \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_N \end{bmatrix} = \begin{bmatrix} \text{IC} \\ \Delta t \mathbf{f}_1 \\ \Delta t \mathbf{f}_2 \\ \vdots \\ \Delta t \mathbf{f}_N \end{bmatrix}$$

where  $\mathbf{I}$  is an identity matrix (of size  $k$ ) and IC are the imposed initial conditions. This system solves for  $N$  timesteps at once with a total number of unknowns  $N \times k$ .

**Remark 1.** *By treating the (unknown) nodal values at different time steps as multiple degrees of freedom associated with each spatial node, we can leverage standard algorithmic approaches (assembly, memory usage) tailored for multiple d.o.f problems. Many approaches in uncertainty quantification (polynomial chaos representation, spectral stochastic methods) leverage such an approach of representing field variation along additional dimensions (stochastic dimensions) as simply additional degrees of freedom at each spatial location [20, 27, 11, 23]. Figure 2(left) illustrates an elemental matrix considering  $N(=10)$  time steps (i.e. 10 degrees of freedom per node), for a 2D quadrilateral element<sup>2</sup>. Note the sparse structure of the elemental matrix as well as the resulting global matrix (for a  $2 \times 2$  quad mesh) in Figure 2 (right).*




FIGURE 2. Left, Nonzero entries in an elemental matrix. Right, Nonzero entries in the resultant global matrix for a 2x2 mesh of quadrilateral elements. Dark squares represent nonzero entries and number are node indices. The time block consists of 10 time steps.

<sup>2</sup>using Backward-Euler discretization in time

**2.2. Space-time framework for non-linear problems.** Extending the approach to certain non-linear problems is straightforward. Consider the case where  $\kappa$  is a function of the dependent variable,  $u$ . We assume that  $\kappa(u)$  satisfies appropriate smoothness and boundedness assumptions to ensure existence and uniqueness,  $0 < \underline{\kappa} \leq \kappa(u) \leq \bar{\kappa} < \infty$ . In this case the weak form for a block  $B$  is

$$(7) \quad (w^h, \frac{u_{n+1}^h - u_n^h}{\Delta t}) + (\nabla w^h, \kappa(u_{n+1}^h) \nabla u_{n+1}^h) = (w^h, f_{n+1}), \text{ for } n = 0, 1, \dots, N$$

The solution to such non-linear equations are usually via (quasi-)Newton schemes. The methodology involves construction of the Jacobian and residual, which are used to compute updates. This is represented in matrix-vector terms as:

$$(8) \quad \mathbf{J}_{\mathbf{u}_{n+1}^i} \delta \mathbf{u}_{n+1}^{i+1} = \mathbf{F}_{\mathbf{u}_{n+1}^i}, \quad \mathbf{u}_{n+1}^{i+1} = \mathbf{u}_{n+1}^i + \delta \mathbf{u}_{n+1}^{i+1},$$

for  $i = 1, \dots$ , until convergence and for  $n = 0, 1, \dots$

where  $\mathbf{J}_{\mathbf{u}_{n+1}^i}$  is the Jacobian (or linearized form), and  $\mathbf{F}_{\mathbf{u}_{n+1}^i}$  is the residual of the above equation, both computed using  $\mathbf{u}_{n+1}^i$ . More specifically, for the non-linear diffusion equations defined by (7), the residual,  $\mathbf{F}_{\mathbf{u}_{n+1}^i}$ , is given by

$$(9) \quad \mathbf{F}_{\mathbf{u}_{n+1}^i} = \frac{1}{\Delta t} \mathbf{M} \mathbf{u}_{n+1}^i - \frac{1}{\Delta t} \mathbf{M} \mathbf{u}_n - \mathbf{K}(\mathbf{u}_{n+1}^i) \mathbf{u}_{n+1}^i - \mathbf{f}_{n+1}$$

where  $\mathbf{K}(\mathbf{u}_{n+1}^i)$  denotes the solution dependent stiffness matrix. The Jacobian is given as:

$$(10) \quad \mathbf{J}_{\mathbf{u}_{n+1}^i} = \frac{1}{\Delta t} \mathbf{M} + \left( \mathbf{K}(\mathbf{u}_{n+1}^i) + \frac{d\mathbf{K}(\mathbf{u}_{n+1}^i)}{du} \right)$$

Instead of sequentially solving for each time step (as in (9)), consider solving for the field variable in a complete time block,  $B$ , consisting of  $N$  timesteps simultaneously. That is,

$$(11) \quad \begin{bmatrix} \mathbf{J}_{\mathbf{u}_1^i} & & & \\ -\frac{1}{\Delta t} \mathbf{M} & \mathbf{J}_{\mathbf{u}_2^i} & & \\ & \ddots & \ddots & \\ & & -\frac{1}{\Delta t} \mathbf{M} & \mathbf{J}_{\mathbf{u}_N^i} \end{bmatrix} \begin{bmatrix} \delta \mathbf{u}_1^{i+1} \\ \delta \mathbf{u}_2^{i+1} \\ \vdots \\ \delta \mathbf{u}_N^{i+1} \end{bmatrix} = \begin{bmatrix} \mathbf{F}_{\mathbf{u}_1^i} \\ \mathbf{F}_{\mathbf{u}_2^i} \\ \vdots \\ \mathbf{F}_{\mathbf{u}_N^i} \end{bmatrix}$$

and

$$\begin{bmatrix} \mathbf{u}_1^{i+1} \\ \mathbf{u}_2^{i+1} \\ \vdots \\ \mathbf{u}_N^{i+1} \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1^i \\ \mathbf{u}_2^i \\ \vdots \\ \mathbf{u}_N^i \end{bmatrix} + \begin{bmatrix} \delta \mathbf{u}_1^{i+1} \\ \delta \mathbf{u}_2^{i+1} \\ \vdots \\ \delta \mathbf{u}_N^{i+1} \end{bmatrix}$$

**Remark 2.** One can alternatively ignore the off-diagonal entries of the block Jacobian to construct an approximate diagonal Jacobian. The propagation of time information is then limited to the residual on the right hand side. We tested both approaches, with the latter approach taking marginally more iterations to convergence, while providing substantial ease of implementation. Unless otherwise stated, all our results are based on the latter approach.

### 3. SPACE-TIME FORMULATION: HIGHER ORDER TIME SCHEMES

We next look at extending the space-time strategy to incorporate two families of higher order time-steppers:  $\Theta$  scheme and Backward difference formula. We consider linear and nonlinear diffusion, and moreover, we also consider the treatment of the Allen-Cahn equation, which is a parabolic PDE with a lower-order nonlinearity, whose solution has evolving layers and for which adaptivity is particularly useful.

**3.1.  $\Theta$  scheme: Linear equation.** The semi-discrete form of the  $\Theta$ -scheme – which is a generalization of the Euler Backward scheme – is as follows:

$$(12) \quad (w, u_{n+1}) - (w, u_n) + \Delta t [(1 - \Theta)(\nabla w, \kappa \nabla u_n) + \Theta(\nabla w, \kappa \nabla u_{n+1})] = \Delta t [(1 - \Theta)(w, f_n) + \Theta(w, f_{n+1})], \text{ for } n = 0, 1, \dots$$

The fully discrete matrix-vector representation is given by

$$(13) \quad \mathbf{M}\mathbf{u}_{n+1} - \mathbf{M}\mathbf{u}_n + \Delta t(1 - \Theta)\mathbf{K}\mathbf{u}_n + \Delta t\Theta\mathbf{K}\mathbf{u}_{n+1} = \Delta t(1 - \Theta)\mathbf{f}_n + \Delta t\Theta\mathbf{f}_{n+1}, \text{ for } n = 0, 1, \dots$$

Again, it is straightforward to group and simultaneously solve for  $N$  time steps together. The corresponding matrix form is expressed as:

$$(14) \quad \begin{bmatrix} \mathbf{I} & & & & \\ -\mathbf{M} + \Delta t(1 - \Theta)\mathbf{K} & \mathbf{M} + \Delta t\Theta\mathbf{K} & & & \\ & -\mathbf{M} + \Delta t(1 - \Theta)\mathbf{K} & \mathbf{M} + \Delta t\Theta\mathbf{K} & & \\ & & \ddots & \ddots & \\ & & & -\mathbf{M} + \Delta t(1 - \Theta)\mathbf{K} & \mathbf{M} + \Delta t\Theta\mathbf{K} \end{bmatrix} \begin{bmatrix} \mathbf{u}_0 \\ \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_N \end{bmatrix} = \begin{bmatrix} \text{IC} \\ \Delta t[(1 - \Theta)\mathbf{f}_0 + \Theta\mathbf{f}_1] \\ \Delta t[(1 - \Theta)\mathbf{f}_1 + \Theta\mathbf{f}_2] \\ \vdots \\ \Delta t[(1 - \Theta)\mathbf{f}_{N-1} + \Theta\mathbf{f}_N] \end{bmatrix}$$

Again, the global space-time matrix (14) has a block structure.

**3.2.  $\Theta$  scheme: Nonlinear diffusion with variable coefficient.** The corresponding weak form for this case is given as:

$$(15) \quad (w^h, \frac{u_{n+1}^h - u_n^h}{\Delta t}) + (1 - \Theta)(\nabla w^h, \kappa(u_n^h) \nabla u_n^h) + \Theta(\nabla w^h, \kappa(u_{n+1}^h) \nabla u_{n+1}^h) = (1 - \Theta)(w^h, f_n) + \Theta(w^h, f_{n+1}), \text{ for } n = 0, 1, \dots$$

The Jacobian is:

$$(16) \quad \mathbf{J}_{\mathbf{u}_{n+1}^i} = \mathbf{M} + \Delta t\Theta \left( \mathbf{K}(\mathbf{u}_{n+1}^i) + \frac{d\mathbf{K}(\mathbf{u}_{n+1}^i)}{du} \right)$$

while the residual is given by:

$$(17) \quad \mathbf{F}_{\mathbf{u}_{n+1}^i} = \mathbf{M}\mathbf{u}_{n+1}^i - \mathbf{M}\mathbf{u}_n + \Delta t\Theta (\mathbf{K}(\mathbf{u}_{n+1}^i)\mathbf{u}_{n+1}^i - \mathbf{f}_{n+1}) + \Delta t(1 - \Theta) (\mathbf{K}(\mathbf{u}_n)\mathbf{u}_n - \mathbf{f}_n)$$

Again, it is straightforward to group and simultaneously solve for  $N$  time steps together. The corresponding (diagonal) block Jacobian (see Remark 2) is given as:

$$(18) \quad \begin{bmatrix} \mathbf{M} + \Delta t \Theta \left( \mathbf{K}(\mathbf{u}_1^i) + \frac{d\mathbf{K}(\mathbf{u}_1^i)}{du} \right) & & & \\ & \mathbf{M} + \Delta t \Theta \left( \mathbf{K}(\mathbf{u}_2^i) + \frac{d\mathbf{K}(\mathbf{u}_2^i)}{du} \right) & & \\ & & \ddots & \\ & & & \mathbf{M} + \Delta t \Theta \left( \mathbf{K}(\mathbf{u}_N^i) + \frac{d\mathbf{K}(\mathbf{u}_N^i)}{du} \right) \end{bmatrix}$$

As before, the upper index denotes Newton-Raphson iteration.

**3.3.  $\Theta$  scheme: Allen Cahn Equation.** The Allen-Cahn equation is a semi-linear diffusion equation with a non-linear reaction term:

$$(19) \quad \partial_t u(\mathbf{x}, t) - \Delta u(\mathbf{x}, t) + \epsilon^{-2} f(u) = 0$$

where we set  $f(u) = u(u^2 - 1)$ . The initial condition is  $u(\mathbf{x}, 0) = u_0$  along with zero flux conditions in the boundaries.

The corresponding semi-discrete form is given as:

$$(20) \quad (w^h, \partial_t u^h) + (\nabla w^h, \nabla u^h) + \epsilon^{-2} (w^h, f(u^h)) = 0$$

Using the  $\Theta$ -scheme results in the fully discrete form:

$$(21) \quad (w^h, u_{n+1}^h) + \Delta t \Theta [(\nabla w^h, \nabla u_{n+1}^h) + \epsilon^{-2} (w^h, f(u_{n+1}^h))] = (w^h, u_n^h) - \Delta t (1 - \Theta) [(\nabla w^h, \nabla u_n^h) + \epsilon^{-2} (w^h, f(u_n^h))]$$

Again, it is straightforward to group and simultaneously solve for  $N$  time steps together. The corresponding (diagonal) block Jacobian is given as:

$$(22) \quad \begin{bmatrix} \mathbf{M} + \Delta t \Theta \left( \mathbf{K}(\mathbf{u}_1^i) + \epsilon^{-2} \frac{d\mathbf{f}(\mathbf{u}_1^i)}{du} \right) & & & \\ & \mathbf{M} + \Delta t \Theta \left( \mathbf{K}(\mathbf{u}_2^i) + \epsilon^{-2} \frac{d\mathbf{f}(\mathbf{u}_2^i)}{du} \right) & & \\ & & \ddots & \\ & & & \mathbf{M} + \Delta t \Theta \left( \mathbf{K}(\mathbf{u}_N^i) + \epsilon^{-2} \frac{d\mathbf{f}(\mathbf{u}_N^i)}{du} \right) \end{bmatrix}$$

Alternative schemes for the Allen-Cahn equation and other phase-field models are described in, e.g., [13].

**3.4. Backward difference formula based time steppers.** BDF based time-steppers of order  $s$  utilize the solution at  $s$  previous time steps to construct the solution at the next time step. A general  $s$  order BDF scheme is given as:

$$(23) \quad \sum_{k=0}^s \alpha_k u_{n+k} = \Delta t \beta g_{n+s}$$

where the left hand side is the BDF scheme representation of the time derivative,  $\frac{\partial u}{\partial t}$ , in terms of the solution  $u_i$  at time point  $i$ , and the right hand side collects all other terms. Here,  $\alpha$  and  $\beta$  are known BDF coefficients [8]. A first order ( $s = 1$ ) BDF scheme is identical to the Euler Backward scheme described earlier. The simplest multi-step scheme is for  $s = 2$  and for the linear diffusion equation it is given as:

$$(24) \quad (w^h, u_{n+2}^h) - \frac{4}{3}(w^h, u_{n+1}^h) + \frac{1}{3}(w^h, u_n^h) + \frac{2}{3}\Delta t(\nabla w^h, \kappa \nabla u_{n+2}^h) = \frac{2}{3}\Delta t(w^h, f_{n+2})$$

Again, it is straightforward to group and simultaneously solve for  $N$  time steps together.<sup>3</sup> The corresponding space-time block equations are as follows:

$$(25) \quad \begin{bmatrix} \mathbf{I} & & & & \\ -\mathbf{M} & \mathbf{M} + \Delta t \mathbf{K} & & & \\ \frac{1}{3}\mathbf{M} & -\frac{4}{3}\mathbf{M} & \mathbf{M} + \frac{2}{3}\Delta t \mathbf{K} & & \\ & & & \ddots & \\ & & \frac{1}{3}\mathbf{M} & -\frac{4}{3}\mathbf{M} & \mathbf{M} + \frac{2}{3}\Delta t \mathbf{K} \end{bmatrix} \begin{bmatrix} \mathbf{u}_0 \\ \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_N \end{bmatrix} = \begin{bmatrix} \text{IC} \\ \Delta t \mathbf{f}_1 \\ \frac{2}{3}\Delta t \mathbf{f}_2 \\ \vdots \\ \frac{2}{3}\Delta t \mathbf{f}_N \end{bmatrix}$$

It is clear that higher order multistep methods produce block matrices that have a larger bandwidth (see Figure 3). Nonlinear problems can be treated similarly.




FIGURE 3. View of nonzero entries of an elemental matrix (2D quad with linear basis functions) for a block of space-time consisting of 2 timesteps + initial condition (left) and 9 timesteps + initial condition(right) with Backward difference formula of second order. First timestep is calculated using Euler Backward scheme.

#### 4. ADAPTIVE MESHING FOR THE BLOCK-SPACE-TIME METHOD: RESIDUAL BASED ERROR ESTIMATOR

A central idea of this work is to develop a block space-time methodology that can be integrated with mesh adaptivity. This will enable targeted refinement of regions that exhibit variations in the corresponding block of time. Mesh adaptivity requires the definition of an indicator function that determines which regions of the space require refinement/coarsening. In this work, we build on prior work and utilize

<sup>3</sup>Note that the first time step is approximated using a backward Euler time stepper as the second order BDF scheme requires knowledge of the solution at two previous time steps



standard residual-based error indicators; see, e.g. [4, 26, 25]. Alternative duality-based indicators for nonlinear parabolic problems are described in, e.g. [9, 10, 12].

Residual based error indicators  $\eta_e$  are constructed for each spatial element,  $\Omega_e$ , and consist of two terms – an interior residual,  $r_{int}$ , and a jump residual,  $r_{jump}$  [26]:

$$(26) \quad \eta_e^2 = h_e^2 \|r_{int}^h\|_{L^2(\Omega_e)}^2 + h_e \|r_{jump}^h\|_{L^2(\partial\Omega_e)}^2$$

where  $h_e$  is the size of element,  $\Omega_e$ . For example, for the linear and quasi-linear diffusion equation, the detailed derivation of the interior and jump residual is available in the work of Verfurth [26]. We refer the interested reader to that work and only show the key results here. Basically, the two terms are constructed (as the name suggests) from the definition of the residual:

$$(27) \quad \mathcal{R}^h(w) = (w, f) - (w, \partial_t u^h) - (\nabla w, \kappa(u^h) \nabla u^h)$$

This residual is decomposed into element-wise terms as:

$$(28) \quad R^h(w) = \sum_e \left( \int_{\Omega_e} w r_{int}^h d\Omega + \int_{\Gamma_e} w r_{jump}^h d\Gamma \right)$$

where

$$(29) \quad r_{int}^h = f - \partial_t u^h + \nabla \cdot \kappa(u^h) \nabla u^h$$

and

$$(30) \quad r_{jump}^h = \begin{cases} 0 & \text{on } \partial\Omega_e \cap \Gamma_h \\ \kappa(u^h) \nabla u_+^h \cdot \mathbf{n}_+ + \kappa(u^h) \nabla u_-^h \cdot \mathbf{n}_- & \text{on } \partial\Omega_e \setminus \Gamma_h \end{cases}$$

where  $\Gamma_h$  is that part of the boundary with Dirichlet conditions imposed.

The interior and jump residuals for the Allen-Cahn equations are similarly defined as

$$(31) \quad r_{int}^h = \epsilon^{-2} f(u^h) - \partial_t u^h + \Delta u^h$$

and

$$(32) \quad r_{jump}^h = \begin{cases} 0 & \text{on } \partial\Omega_e \cap \Gamma_h \\ \nabla u_+^h \cdot \mathbf{n}_+ + \nabla u_-^h \cdot \mathbf{n}_- & \text{on } \partial\Omega_e \setminus \Gamma_h \end{cases}$$

Recall that we are using error indicators defined over a block of time. We extend the concept of error indicator defined at one time step to the notion of an error indicator defined over a block of time steps. There are several choice of error indicators with two of them being:

the average value of the error indicator across the block of time

$$(33) \quad \eta_{e,avg} = \sum_{\text{Block of time, B}} \left( \frac{\sum_{n=1}^N \eta_{e,n}}{N} \right)$$

and the maximum value in the block of time:

$$(34) \quad \eta_{e,max} = \max_{\text{Block of time, B}} \eta_{e,n}$$

The former approach is advantageous for slow variations in the block and avoids frequent migrating of elements between processors, which can negatively affect scalability. The latter approach is advantageous when there is rapid changes across a few timesteps. In this case the error estimator is not 'diffused' by timesteps where solution is changing slowly.

**Remark 3.** *In this work, we approximate the semi-discrete term,  $\frac{\partial u}{\partial t}$  in terms of its finite difference representation ( $\Theta$  or BDF scheme). An implicit assumption is that the time steps are small enough that this approximation is valid. Ideally, one would choose a consistent representation in both space and time; i.e. using a finite element representation for time variations [21, 15]. This provides several advantages in terms of mathematical elegance. We defer this development to a subsequent paper.*

## 5. IMPLEMENTATION DETAILS

We utilize our in-house scalable, parallel FEM framework that is optimized for distributed memory computing. The FEM software library is implemented in C++ and uses object oriented software principles. Linear algebra, parallel matrix and vector storage are all performed by the PETSc library [5]. PETSc modules (KSP, SNES) are used to solve (non)linear equations. The FEM library is dynamically linked to the Parallel Hierarchical Grid (PHG) library [2] which is a parallel mesh refinement framework. PHG uses a bisection type algorithm [29], specifically newest vertex bisection to refine/coarsen elements<sup>4</sup>. PHG operates on simplex elements and produces conforming meshes after refinement. While our existing implementation was reasonably optimized, several software engineering principles had to be carefully implemented to ensure efficient execution of the block space-time problems. Some of the simpler changes that had substantial impact for the space-time formulation<sup>5</sup>, *but had minimal impact on the existing, standard iterative formulation*<sup>6</sup> are: (a) moving calculations out of loops whenever possible, (b) executing calculations and storing results in array before loops, (c) conversion to 64-bit based integer variables for storing matrix indices, which allows going beyond 4 billion unknowns, (d) using local values instead of values obtained through pointer or reference<sup>7</sup>.

We made software engineering decisions to ensure that the space-time implementation is compatible with our existing sequential FEM frameworks (see remark 1). Basically, the space-time formulation is treated as the solution of a steady-state problem. Initial conditions are imposed as boundary condition on the first degree of freedom. We discuss key aspects of the implementation next.

**5.1. Memory interlacing and matrix bandwidth.** We rearrange the vector of unknowns  $\mathbf{u}$  to enumerate time points before looping over space. i.e.  $\mathbf{u} = \{u_1^1, u_2^1, \dots, u_N^1, u_1^2, u_2^2, \dots, u_N^2, \dots, u_N^k\}$ , where subscript refers to time and superscript denotes space. This allows for more efficient assembly, because all data (coefficients in system of equation) for a specific element share memory locality, thus preventing cache misses. Moreover, this approach is compatible with existing frameworks for FEM, because problems with multiple DOFs are supported in existing FEM frameworks, so we can use standard procedures for system assembling or imposing boundary conditions. This has the additional advantage of reducing the matrix

<sup>4</sup>In newest vertex bisection, the edge that lies in opposite to the newest node is divided

<sup>5</sup>even very simple changes possibly like avoiding division operations by replacing with multiplication, or using 'const' to enable the compiler to automatically simplify expressions, and avoiding unnecessary conversions between integer and double values have some impact

<sup>6</sup>these standard best practices provide minimal improvements to sequential time stepping, and are usually not reported or ignored.

<sup>7</sup>the function can access values more quickly, as it does not have to first fetch pointer, than check where pointer points and get value

bandwidth. Table 1 enumerates the bandwidth for space-time formulation (2D mesh with  $100 \times 100$  quad elements, linear basis function) using a Euler Backward formulation. As expected, the bandwidth linearly increases with increasing number of timesteps.

TABLE 1. The bandwidth size for space-time formulation

	number of timesteps				
	sequential	5	10	25	50
row size	10201	51005	102010	255025	510050
bandwidth	205	1023	2043	5103	10203




FIGURE 4. Computational time for different sizes of block size in block based matrix  $aij$  format. Investigated example was the 2D linear diffusion problem using  $100 \times 100$  quad elements with one block of 100 timesteps. Numbers in legend describes block size, the ' $aij$ ' is sparse matrix without blocks and 'no zeroes' is special case with ignored zero entries in element matrix.

**5.2. Matrix access and storage.** We explored the use of block structured matrix formats that are optimized for storing the global matrix, where the block size is a function of the number of timesteps. We use blocks of sizes ranging from  $2 \times 2$  to  $N \times N$  (where  $N$  is the number of timesteps in a block,  $B$ ). Figure 4 plots the time required to assemble and solve a linear diffusion problem using the space-time formulation (Backward Euler). We can see that different non-zero blocks sizes can give observable savings in computational time. Note also that each of these blocks is itself quite sparse (see Figure. 6). Careful identification of the non-zero pattern could potentially result in much larger savings. This is illustrated in the 'no-zero' column in Figure 4 where we avoid inserting zero elements to the global matrix. The memory requirement for this case is also substantially minimized (see figure 5).




FIGURE 5. Memory requirements for different sizes of block size in block based matrix *aij* format and sparse element matrix. On the example of elliptic problem on grid of 100x100 elements with one block of 100 timesteps




FIGURE 6. View of full element matrix for Euler Backward (left) and compressed element matrix (right). Dark grey color means nonzero entries in matrix.

**5.3. Compressed element matrix.** Figure 6 shows the non zero patterns within an elemental stiffness matrix. The size of this matrix is  $(nb f \times N) \times (nb f \times N)$ , where  $nb f$  is the number of basis functions (assuming 1 dof) and  $N$  is the number of time steps in the block,  $B$ . Clearly, as  $N$  increases, the sparsity of the elemental stiffness matrix improves. This suggests using compressed formats for storing the elemental stiffness matrices. This issue becomes more important with increasing block size.

**5.4. Scaling.** We performed scaling studies on two machines. Preliminary scaling was performed on the TACC Stampede [3]. Stampede consists of 6400 compute nodes each equipped with two Intel E5-2680 8-core processors and 32 GB of memory. Each compute node has access to 250GB of local storage and Lustre filesystem with a write performance of 150GB/s. The nodes are connected with FDR InfiniBand

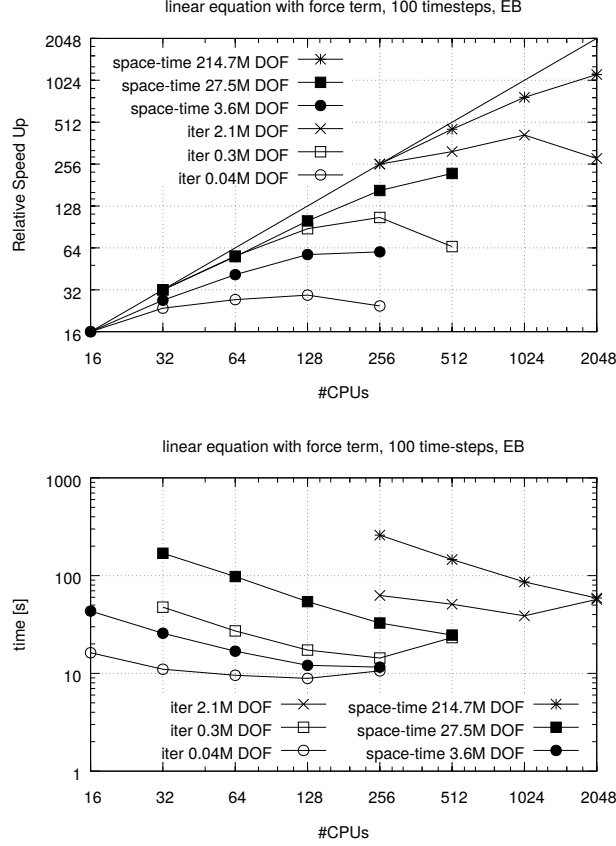


FIGURE 7. Speed up on Stampede (top) and time requirements for space-time and iterative formulations on the same machine (bottom)

network. Scaling results are shown in Figure 7. We also performed scaling on NCSA Blue Waters [1]. Blue Waters consists of 22,640 nodes, each consisting of two AMD 6276 "Interlagos" processors for total number of 362,240 computing cores. Each node has 64GB of memory, which gives total system memory equal to 1.476PB. Compute nodes have access to storage with Lustre filesystem that has size 26.4 PB and offers total bandwidth of more than 1TB/s.

Figures 8, 9, 10 show scaling for the linear, nonlinear and Allen-Cahn equations respectively. We show good scaling upto 131,072 processors. The largest problem size used was  $\sim 5.2$  billion degrees of freedom. This is a  $800^3$  mesh with 10 timesteps per block. The results illustrate reasonably good scaling for the space-time approach, and Figure 7 suggests that for larger number of processors, the space-time approach performs better than the iterative approach. Specifically, beyond a threshold number of processors, the *total time to solve* a given problem is





FIGURE 8. Speed up on Blue Waters for linear time dependent diffusion

FIGURE 9. Speed up on Blue Waters for nonlinear time dependent diffusion (nonlinearity induced by diffusion coefficient depending on  $u$ )

lower for the space time approach compared to the iterative approach. We anticipate that for more complex problems (complex geometries, remeshing, I/O) the space-time approach may perform even better.

## 6. NUMERICAL EXAMPLES

We illustrate our adaptive parallel-in-space-time framework on linear and nonlinear diffusion equation and the Allen-Cahn equation.

**6.1. Problem A – Linear Diffusion.** We first consider the linear case. We set  $\kappa = 1$ . We use the method of manufactured solutions to construct the forcing term in (1) to ensure an analytical solution,  $u$ :

$$(35) \quad u(x, y, z, t) = \exp(-\alpha(x, y, z, t))$$




FIGURE 10. Speed up on Blue Waters Supercomputer for Allen-Cahn problem

where  $\alpha(x, y, z, t)$  is equal to:

$$(36) \quad \alpha(x, y, z, t) = \frac{(x - x_0(t))^2 + (y - y_0(t))^2 + (z - z_0(t))^2}{d^2}$$

and

$$\begin{aligned} x_0(t) &= a \cos(\omega t) + b \\ y_0(t) &= a \sin(\omega t) + b \\ z_0(t) &= b \end{aligned}$$

with  $a = 0.2$ ,  $b = 0.5$ ,  $\omega = 2\pi$ ,  $d = 0.1$ . Thus,  $u$  is a rotating exponential hill centered at the midplane and rotating with a radius of 0.2 and an angular speed of  $2\pi$ . This manufactured solution is constructed by the following forcing term

$$(37) \quad f(\mathbf{x}, t) = - \left( \frac{4((x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2)}{d^4} - \frac{6}{d^2} - 2a\omega \left( \frac{\cos(t\omega)(y - y_0) - \sin(t\omega)(x - x_0)}{d^2} \right) \right) \exp(-\alpha(x, y, z, t))$$

The equation is solved in a region of dimensions  $[0 : 1] \times [0 : 1] \times [0 : 1]$ . Every boundary face of the region has an essential boundary condition with prescribed value of  $u$  equal to the value computed from (35). We use a time step of 0.01 and solve for a block of 100 timesteps.

Figure 11 shows a cutoff of the refined spatial mesh after 20 refinement iterations. We remind the reader that each iteration consists of solving the space-time problem, constructing the time-averaged elemental refinement indicators, equation (33), and then refining the 3D mesh according to the indicators. Given that this is a moving source problem, it is clearly seen that there is refinement in regions where the source has passed through.

We compare spatial convergence rates for three implementations: (a) sequential time stepping with no spatial adaptivity, (b) space-time implementation with no spatial adaptivity, and (c) space-time implementation with spatial adaptivity.






FIGURE 11. Refined mesh: view from top (left) and front (right).

FIGURE 12. Grid convergence for linear heat equation. Error is  $\|u - u_h\|_2$  from last time-step versus average element size  $\bar{h}$ .FIGURE 13. Time convergence for linear heat equation. Error is  $\|u - u_h\|_2$  from last time-step versus size of time-step.

We plot convergence in Figure 12, where error is  $\|u - u_h\|_2$ . The first two implementations (obviously) overlap, with the adaptive mesh implementation showing a



reduced error. All three curves show a slope of 2, which is to be expected. The Figure 13 shows time-step convergence, with expected slopes of 1 and 2 for Backward Euler and Backward Difference Formulae, respectively.

**6.2. Problem B – Nonlinear Diffusion.** For the nonlinear case, we set the coefficient  $\kappa(u) = 1 + 10u^2$  and we choose an exact solution such that  $|u| \leq 1$ , thus bounding  $\kappa$ . As before, we choose our analytical solution to be given by (35). The forcing term,  $f$ , is consequently:

$$(38) \quad f(\mathbf{x}, t) = \frac{80u}{d^4} ((x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2) \exp(-2\alpha(x, y, z, t)) + (1 + 10u^2) \left( \frac{4((x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2)}{d^4} - \frac{6}{d^2} \right) \exp(-\alpha(x, y, z, t)) - 2a\omega \frac{\cos(t\omega)(y - y_0) - \sin(t\omega)(x - x_0)}{d^2} \exp(-\alpha(x, y, z, t))$$

Figure. 14 plots several time snapshots of the moving non-linear source problem, with the solution accurately tracking the moving source. In Figure. 15, we plot spatial convergence for three implementations: (a) sequential time stepping with no spatial adaptivity, (b) space-time implementation with no spatial adaptivity, and (c) space-time implementation with spatial adaptivity (plotted with the average element size). Convergence rates follow along expected lines, with a slope of 2.

**6.3. Problem C – Allen-Cahn.** In this final example we solve the modified Allen-Cahn problem. The key physics (i.e. phase change) which is described by this non-linear equation essentially occurs in a highly localized region of the domain (on a surface of co-dimension 1). Adaptive refinement (and coarsening) has been a very effective approach to accurately resolve this localized region. The equation is given as  $\frac{\partial u}{\partial t} = -D(f(u) - C_n^2 \nabla^2 u)$ , where

$$(39) \quad f(u) = 2Au(1 - 3u + 2u^2) - k$$

and  $D = 1$ ,  $C_n = 0.1$ ,  $A = 16$ ,  $k = 0.1$ . The initial conditions are

$$(40) \quad u(x) = 0.5 + 0.5 \tanh \left( \frac{r - 0.5}{\sqrt{\frac{2}{A} C_n}} \right), \quad r = \sqrt{x^2 + y^2 + z^2}$$

with zero flux conditions on all boundaries. This represents an initial solid of radius,  $r = 0.5$ , that is melting. Using symmetry arguments, we consider a single octant of the space  $([0 : 1] \times [0 : 1] \times [0 : 1])$ . We consider a time step of 0.02 and a block size of 50 time-steps. Figure 16 shows snapshots of the melting sphere at various time points.

Figure 17 illustrates adaptive mesh refinement across the time block. The figure shows the refined mesh for the first block of time ( $t = 0$  to  $t = 1$ ) Note that refinement is localized along the solidification front *within the time block*. To capture this thin interface using a uniform mesh would have required 589 824 elements, instead of the 137 776 elements that were used here.




FIGURE 14. Results of computations where upper half of computational domain is removed. Values of  $u$  change in time from 0s (a) to 1s (d)

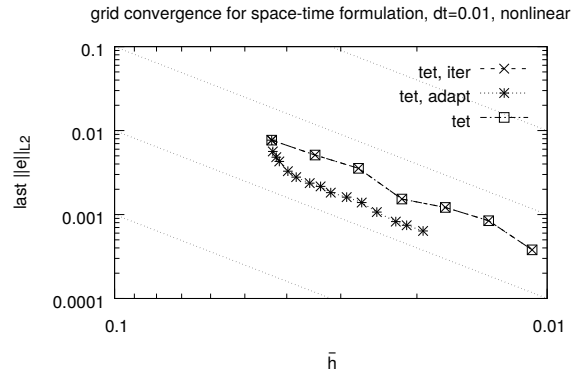


FIGURE 15. Grid convergence for nonlinear heat equation. Error is  $\|u - u_h\|_2$  from last time-step versus average element size  $\bar{h}$ .




FIGURE 16. Results of computations for the 3D Allen-Cahn problem. Values of  $u$  for initial condition (a), after 1s (b), 3s (c) and 5s (d)




FIGURE 17. Finite element mesh after 20 refinement iterations with superimposed solution at time 1s at the dense part of mesh

## 7. CONCLUSION

We present formulation, implementation details and representative examples of a parallel-in-space-time based adaptive methodology for the solution of (linear and) non-linear time dependent problems. The basic concept is to solve for large blocks of space-time unknowns instead of marching sequentially in time. The methodology is a combination of a computationally efficient implementation of a parallel-in-space-time finite element solver coupled with *a posteriori* space-time error estimates and a

parallel mesh generator. We illustrate how this implementation is especially tailored for massively parallel computations. We show good scaling behavior up to 150,000 processors on the Blue Waters machine. This methodology enables scaling on next generation multi-core machines by simultaneously solving for large number of time-steps, and reduces computational overhead by refining spatial blocks that can track localized features. This methodology also opens up the possibility of efficiently incorporating adjoint equations for error estimators and inverse design problems, since blocks of space-time are simultaneously solved and stored in memory. Our future work is focused on extending the space-time framework to utilizing finite element basis functions in time (which enables formal derivation of space-time a posteriori error estimates), and subsequently implementing 4D finite elements to enable simultaneous space and time adaptivity.

## REFERENCES

1. *Blue waters user portal | system summary*.
2. *Phg (parallel hierarchical grid)*.
3. *Texas advanced computing center - stampede technical details*.
4. Mark Ainsworth and J. Tinsley Oden, *A posteriori error estimation in finite element analysis*, John Wiley & Sons, New York, 2000.
5. Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith, *Efficient management of parallelism in object oriented numerical software libraries*, Modern Software Tools in Scientific Computing (E. Arge, A. M. Bruaset, and H. P. Langtangen, eds.), Birkhäuser Press, 1997, pp. 163–202.
6. Marek Behr, *Simplex space-time meshes in finite element simulations*, Int. J. Numer. Meth. Fluids **57** (2008), no. 9, 1421–1434.
7. Kevin Burrage, *10. parallel methods for systems of ordinary differential equations*, Applications on Advanced Architecture Computers, 1996, pp. 101–120.
8. John C. Butcher, *Numerical methods for ordinary differential equations*, John Wiley & Sons, West Sussex, 2008.
9. V. Carey, D. Estep, A. Johansson, M. Larson, and S. Tavener, *Blockwise adaptivity for time dependent problems based on coarse scale adjoint solutions*, SIAM J. Sci. Comput. **32** (2010), no. 4, 2121–2145.
10. G. Şimşek, X. Wu, K.G. van der Zee, and E.H. van Brummelen, *Duality-based two-level error estimation for time-dependent pdes: Application to linear and nonlinear parabolic equations*, Computer Meth. Appl. Mech. Engrg. **288** (2015), 83–109.
11. Bert J. Debuschere, Habib N. Najm, Philippe P. Pbay, Omar M. Knio, Roger G. Ghanem, and Olivier P. Le Matre, *Numerical challenges in the use of polynomial chaos representations for stochastic processes*, SIAM J. Sci. Comput. **26** (2004), no. 2, 698–719.
12. K. Eriksson, C. Johnson, and A. Logg, *Adaptive computational methods for parabolic problems*, Encyclopedia of Computational Mechanics, in: Fundamentals (E. Stein, R. de Borst, and T.J.R. Hughes, eds.), John Wiley & Sons, Ltd., 2004, pp. 675–702.
13. Hector Gomez and Kristoffer G. van der Zee, *Computational phase-field modeling*, Encyclopedia of Computational Mechanics, Second Edition, John Wiley & Sons, Ltd., (In Press).
14. Jeffrey Hittinger, Sven Leyffer, and Jack Dongarra, *Models and algorithms for exascale computing pose challenges for applied mathematicians*, SIAM News, 2013.
15. Thomas J. R. Hughes and G. M. Hulbert, *Space-time finite element methods for elastodynamics: Formulations and error estimates*, Comput. Methods Appl. Mech. Eng. **66** (1988), no. 3, 339–363.
16. Thomas J.R. Hughes and James R. Stewart, *A space-time formulation for multiscale phenomena*, J. Comput. Appl. Math. **74** (1996), no. 1, 217–229.
17. J. Lions, Yvon Maday, and Gabriel Turinici, *A “parareal” in time discretization of pde’s*, C. R. Acad. Sci. Series I Mathematics **332** (2001), no. 7, 661–668.

18. Robert B. Lowrie, Philip L. Roe, and Bram van Leer, *Space-time methods for hyperbolic conservation laws*, Barriers and Challenges in Computational Fluid Dynamics (V. Venkatakrishnan, Manuel D. Salas, and Sukumar R. Chakravarthy, eds.), Springer Netherlands, Dordrecht, 1998, pp. 79–98.
19. Karthik Mani and Dimitri Mavriplis, *Efficient solutions of the euler equations in a time-adaptive space-time framework*, 49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, American Institute of Aeronautics and Astronautics, 2011.
20. Akil Narayan and Dongbin Xiu, *Stochastic collocation methods on unstructured grids in high dimensions via interpolation*, SIAM J. Sci. Comput. **34** (2012), no. 3, A1729–A1752.
21. J. P. Pontaza and J. N. Reddy, *Space-time coupled spectral/hp least-squares finite element formulation for the incompressible navier-stokes equations*, J. Comput. Phys. **197** (2004), no. 2, 418–459.
22. Thomas C. S. Rendall, Christian B. Allen, and Edward D. C. Power, *Conservative unsteady aerodynamic simulation of arbitrary boundary motion using structured and unstructured meshes in time*, Int. J. Numer. Meth. Fluids **70** (2012), no. 12, 1518–1542.
23. Christian Soize and Roger Ghanem, *Physical systems with random uncertainties: Chaos representations with arbitrary probability measure*, SIAM J. Sci. Comput. **26** (2004), no. 2, 395–410.
24. Tayfun E. Tezduyar, Sunil Sathe, Ryan Keedy, and Keith Stein, *Spacetime finite element techniques for computation of fluidstructure interactions*, Computer Meth. Appl. Mech. Engrg. **195** (2006), no. 17–18, 2002–2027.
25. Rüdiger Verfürth, *A review of a posteriori error estimation and adaptive mesh-refinement techniques*, Wiley-Teubner, New York and Stuttgart, 1996.
26. ———, *A posteriori error estimation techniques for finite element methods*, Oxford University Press, Oxford, 2008.
27. Xiaoliang Wan and George Em Karniadakis, *Multi-element generalized polynomial chaos for arbitrary probability measures*, SIAM J. Sci. Comput. **28** (2006), no. 3, 901–928.
28. Luming Wang and Per-Olof Persson, *A high-order discontinuous galerkin method with unstructured spacetime meshes for two-dimensional compressible flows on domains with large deformations*, Comput Fluids **118** (2015), 53–68.
29. Lin-Bo Zhang, *A parallel algorithm for adaptive local refinement of tetrahedral meshes using bisection*, Numer. Math. Theor. Meth. Appl. **2** (2009), no. 1, 65–89.